

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED FINAL/30 SEP 92 TO 30 SEP 94
----------------------------------	----------------	--

4. TITLE AND SUBTITLE DESIGNING EFFICIENT MULT-AGENT SYSTEMS	5. FUNDING NUMBERS
--	--------------------

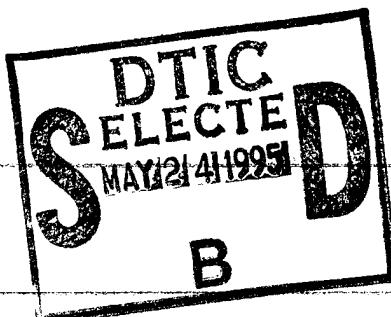
6. AUTHOR(S) YOAV SHOHAM	2304/GS F49620-92-J-0547
------------------------------------	-----------------------------

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) COMPUTER SCIENCE DEPARTMENT STANFORD UNIVERSITY STANFORD, CA 94305	8. PERFORMING ORGANIZATION REPORT NUMBER
---	---

AFSR-IR-95-0351

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) AFOSR/NM 110 DUNCAN AVE, SUITE B115 BOLLING AFB DC 20332-0001	10. SPONSORING/MONITORING AGENCY REPORT NUMBER F49620-93-0547 F49620-92-J-0547
---	--

11. SUPPLEMENTARY NOTES



12a. DISTRIBUTION AVAILABILITY STATEMENT
--

12b. DISTRIBUTION CODE

APPROVED FOR PUBLIC RELEASE: DISTRIBUTION IS UNLIMITED

13. ABSTRACT (Maximum 200 words)

The research carried out during 1994 under this grant expands our previous work on two central design issues in multi-agent systems and begins to develop two new tools for multi-agent systems.

As part of our effort to provide foundations for more powerful agent-oriented programming environments (AOP), we have been concentrating on basic understanding of mental states and their dynamics. In order to facilitate efficient interactions in shared multi-agent environments, such as the Internet, we have continued our investigation of computational social laws and their implementation in different domains. In order to help model agents within a multi-agent systems we have developed an approach for ascribing mental qualities to agents. Finally, we have suggested an approach for describing the complexity of robotics domains, which has yielded an algorithm for analyzing the possibility of decentralizing tasks.

14. SUBJECT TERMS	15. NUMBER OF PAGES
DTIC QUALITY INSPECTED 5	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED
19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT SAR (SAME AS REPORT)

Report on results from research carried out under AFOSR grant
F49620-92-J-0547 during 1994.

Designing Efficient Multi-Agent Systems

by

Yoav Shoham
Computer Science Department
Stanford University
Stanford, CA 94305

1 Overview

The research carried out during 1994 under this grant expands our previous work on two central design issues in multi-agent system and begins to develop two new tools for multi-agent systems.

As part of our effort to provide foundations for more powerful agent-oriented programming environments (AOP)[13], we have been concentrating on basic understanding of mental states and their dynamics. In order to facilitate efficient interactions in shared multi-agent environments, such as the Internet, we have continued our investigation of computational social laws and their implementation in different domains. In order to help model agents within a multi-agent system we have developed an approach for ascribing mental qualities to agents. Finally, we have suggested an approach for describing the complexity of robotics domains, which has yielded an algorithm for analyzing the possibility of decentralizing tasks.

2 Research on AOP and mental states

AOP is a computational framework that can be viewed as a specialization of OOP, where the state of an agent consists of components called beliefs, choices, capabilities, commitments, and possibly others; for this reason the state of an agent is called its *mental state* [14]. The mental state of agents is captured formally in an extension of standard epistemic logics: beside temporalizing the knowledge and belief operators, AOP introduces operators for commitment, choice and capability. Agents are controlled by *agent programs*, which include primitives for communicating with other agents. In the spirit of *speech-act theory*, each communication primitives is of a certain type: informing, requesting, offering, and so on.

A detailed discussion of AOP appears in [13] and in the past we have implemented an agent interpreter: it is documented in [18], and also described in [17]. Past collaboration with the Hewlett Packard corporation involved incorporating features of AOP in the New WaveTM architecture.

Our past work on AOP emphasizes the importance of basic understanding of the notion of a mental state, especially the notions of belief and knowledge. Our results in this area provide for a better understanding of two different approaches to belief revision [5] and a unified approach to belief revision and update [8]. We have also shown a method in which criteria for belief revision can be derived from basic principles [7]. In order to facilitate actual computation with databases of beliefs, we have provided a compilation technique that allows for fast on-line inference [6]. Finally, the work of [10] provides a powerful extension of our previous results [11] on the relation between knowledge and belief.

3 Coordination in open distributed systems

In open multi-agent systems, such as the Internet, different agents are designed by different designers, yet all must share common resources. A major problem is that of coordinating the use of shared resources, such as space,

cpu time, communication channels etc. We refer to such systems as artificial social systems, and have proposed in the past the notion of computational social laws as a design paradigm for such systems [16]. This year we have continued the development of the theory of social laws and have demonstrated the validity of this approach in an important application.

On the theoretical side we have examined the issue of emergent conventions. This work is motivated by the realization that in open distributed systems one cannot foresee all the possible conflicts that can arise ahead of time. In order to overcome this problem we suggest in [9, 15] that agents be provided with the capability to design behavioral conventions on-line, e.g., agents should be able to agree on how to divide CPU time. We have examined a number of approaches for enhancing agents' abilities to form conventions, and have identified one approach that seems particularly promising.

These ideas have been applied to the problem of load balancing [12], a domain in which different agents share computational resources. Our results demonstrate the validity of our approach, which is both adaptive and completely distributed.

4 Modeling agents

For agents to efficiently act within a multi-agent system, some ability to predict the actions of other agents may be required. Predicting the actions of another agent often requires constructing a model of that agent. Low-level models are often difficult to construct, since they require intimate knowledge of the other agent's structure. We have proposed using high level models based on the notion of a mental state, which we call *mental-level* models [3, 1]. These models represent agents *as if* they had a mental state, and they can be constructed based on the agent's past actions together with some high-level background knowledge. Unlike our work on AOP, where the mental states are used to design the agent, here we do not assume that the agents actually have a mental state, but only model them as having one.

5 Analyzing robotics task and their distribution

As part of our ongoing research on mental states, we have investigated the applications of a formal notion of knowledge in robotics. We have been able to use this to obtain new results on motion planning [2]. These results have lead us to propose a more ambitious role for knowledge as an analysis tool for robotics task [4]. We are able to characterize the difficulty of a robotics task by the amount of knowledge required to carry it out, and we characterize the capabilities of a robot by the amount of knowledge it can attain in a given environment. We have used these ideas to provide an algorithm for eliminating the need for a central controller in multi-agent systems. This algorithm analyzes the task at hand and the requirements it makes on each of the agent, and tells us what modifications we must make to each agent so that the task can be performed in a distributed fashion.

6 Presentations

The work described in this report has been presented in a number of conferences, including three presentations in KR'94, two presentations in AAAI'94 and one presentation in TARK'94.

References

- [1] *Action Prediction using a Mental-level model*. 1995.
- [2] R. I. Brafman, J. C. Latombe, Y. Moses, and Y. Shoham. Knowledge as a tool in motion planning under uncertainty. In R. Fagin, editor, *Proc. 5th Conf. on Theor. Asp. of Reas. about Know.*, pages 208–224, San Francisco, 1994. Morgan Kaufmann.

- [3] R. I. Brafman and M. Tennenholtz. Belief ascription and mental-level modelling. In J. Doyle, E. Sandewall, and P. Torasso, editors, *Proc. of Fourth Intl. Conf. on Principles of Knowledge Representation and Reasoning*, pages 87–98, 1994.
- [4] Ronen I. Brafman and Yoav Shoham. Knowledge considerations in robotics and distribution of robotic tasks. submitted to IJCAI'95, 1995.
- [5] Alvaro del Val. On the relation between the coherence and foundations theories of belief revision. In *AAAI'94. Proceedings of the Twelfth National American Conference on Artificial Intelligence*, pages 909–914, 1994.
- [6] Alvaro del Val. Tractable databases: How to make unit resolution complete through compilation. In J. Doyle, E. Sandewall, and P. Torassi, editors, *KR'94, Proceedings of Fourth International Conference on Principles of Knowledge Representation and Reasoning*, pages 551–561, 1994.
- [7] Alvaro del Val and Yoav Shoham. Deriving properties of belief update from theories of action. *Journal of Logic, Language, and Information*, 3:81–119, 1994. Also as Technical Report STAN-CS-TR-94-1501, Computer Science Department, Stanford University.
- [8] Alvaro del Val and Yoav Shoham. A unified view of belief revision and update. *Journal of Logic and Computation*, 4:797–810, 1994. Special Issue on Actions and Processes.
- [9] James E. Kittock. The impact of locality and authority on emergent conventions. In *Proceedings of the Twelfth National Conference on Artificial Intelligence (AAAI '94)*, pages 420–425, 1994.
- [10] P. Lamarre and Yoav Shoham. Knowledge, certainty, belief and conditionalization. In *Proc. of Fourth Intl. Conf. on Principles of Knowledge Representation and Reasoning*. 1994.

- [11] Y. Moses and Y. Shoham. Belief as defeasible knowledge. *Artificial Intelligence*, 64(2):299–322, 1993.
- [12] A. Schaefer, Y. Shoham, and Moshe Tennenholtz. Adaptive load balancing: a study in multi-agent learning. *Journal of AI Research*. To appear.
- [13] Y. Shoham. Agent-oriented programming. *Artificial Intelligence*, 60(1):51–92, 1993.
- [14] Y. Shoham and S. B. Cousins. Logics of mental attitudes in ai. In G. Lakemeyer and B. Mebel, editors. *Advances in Knowledge Representation and Reasoning*. 1994.
- [15] Y. Shoham and M. Tennenholtz. Emergence of social conventions in a game-theoretic setting. Submitted to Artificial Intelligence.
- [16] Y. Shoham and M. Tennenholtz. On social laws for artificial agent societies: Off-line design. *Artificial Intelligence*. To appear.
- [17] Y. Shoham. AGENT0: a simple agent language and its interpreter. In *Proc. NCAI*, Anaheim, CA. 1991.
- [18] M. Torrance. The AGENT-0 programming manual (revise), 1991. Computer Science Department, Stanford University.

Accession For	
NTIS GRA&I <input checked="" type="checkbox"/>	
DTIC TAB <input type="checkbox"/>	
Unannounced <input type="checkbox"/>	
Justification	
By _____	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
R-1	